**CESM Tutorial: Ice- and Ocean-Model Exercises**

The following instructions will make the most sense to people who are familiar with the CESM User's Guide protocols and terminology. Because you are new CESM users, these exercises are fairly self-contained and include most of the information that you will need. However, as the exercises progress, you may find that you need to look up a detail or two in the CESM documentation. We encourage you to refer to the documentation for complete information on how to create, configure, build, and run the following cases and how to customize the CICE and POP2 models. All guides are available on the CESM1.0 website (http://www.cesm.ucar.edu/models/cesm1.0/ )

You may not have time to run all of the following exercises. You should plan to complete the first two cases, then choose either CASE3 or CASE4, and then work on the Data Analysis exercise. If you find yourself blasting through all of these exercises and have time to spare, we have included an Advanced Exercises section at the end. We encourage you to try them out while you have voluteers available to help you.

---

**CASE1: Baseline Case**

---

**I. Using the tutorial version of the CESM, create and configure a low-resolution 1850 ice-ocean ecosystem case on bluefire.** The **G_1850_ECOSYS** compset will produce a case that includes the active ice model (cice), the active ocean model (pop2) with ecosystem model, and data atmosphere and land models using normal-year forcing datasets.

1. The commands in this set of exercises are csh/tcsh commands. Define your shell to be tcsh, or else translate the following commands to your preferred shell.

2. Define the CESM1 tutorial location.
```
> setenv  CCSMROOT /gpfs/proj2/fis/cgd/cseg/csm/collections/cesm1_0_tutorial
```

3. Define your CASE name. *Recommendation: organize your CESM cases in a logical, orderly fashion*
```
> setenv MYRUNS ~/cesm1/tutorial/runs
> setenv  CASE1  $MYRUNS/g_1850_ecosys.base
```

4. Create your case.
```
> cd $CCSMROOT/scripts
>./create_newcase  -compset  G_1850_ECOSYS  -res  T62_gx3v7  -mach  bluefire  -case $CASE1
```

5. Modify your pe-layout for improved throughput by putting the ocean model on its own set of 48 pes. *Recommendations:*
   - *out-of-the-box pe-layouts are not always optimal. You should get into the habit of checking the layout prior to running and be prepared to change the layout to improve efficiency, if needed.*
   - *keep a copy of the original version of any file that you modify; it will help you track your changes*

```
> cd $CASE1
> cp env_mach_pes.xml env_mach_pes.xml.orig
> $EDITOR env_mach_pes.xml
    <entry id="NTASKS_OCN"   value="48"  />
    <entry id="NTHRDS_OCN"   value="1"  />
    <entry id="ROOTPE_OCN"   value="16"  />
```

6. Configure your case.

```
> ./configure -case
```

7. *Recommendation: add notes to the $CASE1/README.case file, to help you document your changes. When you return to your case at a later time, you'll most likely find these notes useful.*

```
> $EDITOR $CASE1/README.case
```

## II. Customize your ice- and ocean-model output.

In this part of the exercise, you will gain experience customizing your ice- and ocean-model output.

1. In the ocean model, customize your case to produce a new (ie, not set up by default) output file containing a daily averaged timeseries of sea-surface heights. To do this, you would normally follow the instructions in the POP User Guide section titled “Time-averaged history files,” but here's a time-saver for you:

```
> $EDITOR $CASE1/Buildconf/pop2.buildnml.csh
```

   Find the namelist tavg_nml (search for the string POP2_TAVG_NML_BASE to get there quickly) and then change or add the following variables to the values shown below in **bold**:

   n_tavg_streams = **4**
   tavg_freq_opt = 'nmonth' 'nday' 'once' **'nday'**
   tavg_freq = 1 1 1 **1**
   tavg_file_freq_opt = 'nmonth' 'nmonth' 'once' **'nyear'**
   tavg_file_freq = 1 1 1 **1**
   tavg_stream_filestrings = 'nmonth1' 'nday1' 'once' **'SSH.nday1'**

```
> cd $CASE1/SourceMods/src.pop2
> cp $CCSMROOT/models/ocn/pop2/input_templates/gx3v7_tavg_contents .
> cp  gx3v7_tavg_contents  gx3v7_tavg_contents.orig
> $EDITOR gx3v7_tavg_contents
```

   change 1 SSH to **4** SSH

2. To turn on high-frequency, daily output in the sea-ice model, we need to modify some namelist variables in $CASE1/Buildconf/cice.buildnml.csh. In the "setup_nml" part, change the histfreq as follows:

   histfreq = 'm','d','x','x','x'

   This leaves the monthly stream as the first history stream and adds a daily history stream to the second stream. Once this is added, we need to tell the ice model which variables will go on the daily output stream. This is done by browsing the "icefields_nml" and modifying the variables as follows:

   f_aisnap = 'mdxxx'
   f_hisnap = 'mdxxx'

   This implies that snapshots of ice concentration and ice thickness will be written to both the monthly and the daily output streams. You will find that these variables are already set to the values you want, so you do not need to change them.

## III. Build and run:

**CESM Tutorial: Ice- and Ocean-Model Exercises**

1. Build $CASE1. *Recommendation: build the case in the background and capture your build output in an output file; check on the progress of the build from time to time.* ***The build will take about six minutes, so you can work ahead on the next exercise, or maybe take a quick break.*** *Re-check the progress of the build after about five to six minutes.*

```
> cd $CASE1
> $CASE1/*.build  >&!  buildout &
> tail buildout
```

2. Sanity check: after a successful build, run the case for five days (default setup). Check on job status.

```
> bsub < *.run
> bjobs
```

3. Confirm that the run was successful by viewing the component model output log files in your $LOGDIR

```
>source $CASE1/Tools/ccsm_getenv
> ls $LOGDIR
```

   If the files from your run are not in the $LOGDIR, then it's likely that your job failed and you should look at the log files in your $RUNDIR directory to find out why.

4. Run the case for one month, with an automatic one-month resubmission.

```
> cd $CASE1
> xmlchange -file env_run.xml -id STOP_OPTION -val nmonths
> xmlchange -file env_run.xml -id STOP_N      -val 1
> xmlchange -file env_run.xml -id RESUBMIT -val 1
>  bsub < *.run
```

5. Review your case's performance and output.
   - Check the load balancing. What would you do to improve the throughput? Why is the present pe layout better than the default? Why did we put the ocean on 48 pes?
   - What ice and ocean output files are available? What are those **.rh.*** files, anyway? Did you get all of the output files that you expected?

6. Questions. *Recommendation: return to this part after completing CASE2, CASE3/4, and data analysis exercises.*
   - When you modified the ocean tavg_nml namelist, you didn't fully specify all of the information for the fourth output stream. Why did the model run correctly anyway?
   - What would you need to do in order to have SSH appear in two output streams (eg, the first stream and the fourth). If you have time, test your method and see what happens.

---

## <u>CASE2</u>: Disable Ocean-Model Tidal Mixing

---

Now that you have run a two-month baseline case, you're ready to set up and run your first scientific experiment. To create this "one-off" experiment from your baseline case, use the handy CESM1 script, **create_clone**. This script will create a "clone" of your baseline case. Repeated use of create_clone will allow you to create a series of one-off scientific experiments.

## CESM Tutorial:  Ice- and Ocean-Model Exercises

***Recommendation****: It is important to note that unless you are an experienced user, you should not change env_conf.xml in your cloned case.  You can do this, but it is beyond the scope of these exercises to walk you through the process.*

### I. Create a clone of $CASE1

1. Define your new case name.  *Recommendation: organize your CESM cases in a logical, orderly fashion. If you are creating a series of experiments, you might want to set up some sort of case-name ordering system, such as adding a suffix of the form ".nn"*

```
> setenv CASE2 $MYRUNS/g_1850_ecosys.01
```

2. Clone $CASE1

```
> cd $CCSMROOT/scripts
> ./create_clone -case $CASE2 -clone $CASE1
> cd $CASE2
> ./configure -case
```

3. ***Recommendation****: add notes to the $CASE2/README.case file. You might add something like this: "This is a clone of  the baseline case, in which the ocean tidal mixing is disabled" Note that the README.case file is an exact copy of the $CASE1/README.case file, which is a bit confusing. You might want to add comments to the $CASE2/README.case file to help clarify.*

```
> $EDITOR $CASE2/README.case
```

4. ***Recommendation****: Notice that all of your customizations to the ocean and ice build scripts in CASE1,* pop2.buildnml.csh and cice.buildnml.csh,  *have been preserved in CASE2.*

### II. Make your science changes in $CASE2

1. Deactivate the ocean-model tidal mixing option**.**  You should check the POP documentation for full information, but for expediency, we'll tell you what you need to do: find the tidal_nml  namelist in $CASE2/Buildconf/pop2.buildnml.csh and set ltidal_mixing false.

2. Build $CASE2. Set CONTINUE_RUN FALSE using the xmlchange command or by editing the env_run.xml file. Run the job for five days. Check job status frequently.

```
> $CASE2/*.build  >&! $CASE2/buildout &
> xmlchange -file env_run.xml -id CONTINUE_RUN -val FALSE
> bsub < *.run
> bjobs
```

3. Check run status.  Something went wrong with this case. What is the problem?  How would you fix it? *Hint: use "ls -lt" to identify the most recent output log files in your run directory. Better yet: try the "dired" utility:  ~njn01/bin/dired -t*

---

### <u>CASE3</u>:   Change Ocean-Model Anisotropic Viscosity Alignment

---

## CESM Tutorial: Ice- and Ocean-Model Exercises

Now that you've experienced a failed science experiment, you're probably ready to set up and run a *successful* one. Choose either this case or CASE4. In either exercise, you will create another "one-off" experiment from your baseline case.

***Remember**: It is important to note that unless you are an experienced user, you should not change env_conf.xml in your cloned case. You can do this, but it is beyond the scope of these exercises to walk you through the process.*

**I. Create a clone of $CASE1.** Define a new case name ($CASE3) and clone $CASE1 as in the previous example.

**II. Make your science changes in $CASE3**

1. Change the ocean-model anisotropic viscosity alignment option. Change the default POP settings to select the alignment option for the parallel component of the anisotropic viscosity. For expediency, here's what you need to do: Find the hmix_aniso_nml namelist in your pop2.buildnml.csh and select the 'flow' option. The details are in the POP User Guide.

2. Build $CASE3 and run it for two months.

3. Compare your output with the baseline case using tools discussed in earlier tutorial sessions. For quick comparisons, you might use the ncdiff and ncview commands.

**<u>CASE4</u>: Change Ice-Model Snow and Sea Ice Albedo**

---

**I. Create a clone of $CASE1.** Define a new case name ($CASE4) and clone $CASE1 as in the previous examples. ***Remember**: It is important to note that unless you are an experienced user, you should not change env_conf.xml in your cloned case. You can do this, but it is beyond the scope of these exercises to walk you through the process.*

**II. Make your science change in $CASE4**

1. Change the "snow and sea ice albedo". Note that we are actually changing a more fundamental optical property of the sea ice and snow. Find the R_snw and R_ice namelist parameters in the CICE users guide. The details of these are in the scientific reference guide, but they are essentially a tuning parameter that specifies the number of standard deviations away from the base optical properties in the shortwave radiative transfer code. The albedo that is computed based on these optical properties is nearly proportional to these parameters.

2. First, what sign would you use to decrease the snow albedo? What sign would you use to increase the snow albedo?

3. Let's hit it pretty hard now. Try decreasing the snow and ice optical properties (R_snw and R_ice) by 3 standard deviations each. Set these in the cice.buildnml.csh file.

4. Build $CASE4 and run it for two months.

5. Compare your output with the baseline case.

6. Keep in mind: What time of year did you start your run? Which season do you expect to see the biggest impact for shortwave changes?

## CESM Tutorial: Ice- and Ocean-Model Exercises

---

## Data Analysis Exercise

---

The baseline case, $CASE1, is not a recommended scientific configuration. If you had enough time today to run this case for a decade and analyze the results, you would see problems in the model solutions. But because you don't, we have done that for you. Data are located on bluefire:/ptmp/njn01/g_1850_ecosys.base.production

Using the techniques you have learned in this workshop, view the model results and describe the problems you see.

Suggestions for further analysis:

- create a 10-year time series of monthly average ocean SSTs (*hint: use ncrcat with the -v option*)
- create decadal averages of ice and/or ocean velocities (*hint: use ncdump to determine the names of the velocity fields and ncra to create the averages*)

---

## Advanced Exercises

---

Here is your opportunity to explore different ice and ocean model options. You might consider the following exercises, or come up with one of your own. Discuss with one of the volunteers if you need help or want some feedback.

*Recommendation: divide up the following exercises within your group, so that together you can cover most or all of these cases and share your results, thereby leveraging your efforts and maximizing your learning.*

1. Set up and run a case in which you double the ice- or ocean-model wind stress. *Hint: modify the component model's _comp_mct.F90 file at the point where the model receives the wind stress from the coupler.*
2. How would you set up a case to change the sea-ice model shortwave physics back to the CCSM3 formulation?
3. Set up and run a case in which you couple the ocean model every hour and activate the constant short-wave distribution option. You'll need to cut the ocean-model timestep; try a factor of two. What would be the purpose of such a run?
4. Set up and run an ocean-only case in which you turn the overflow parameterization off. You might do this to see the effect of overflows on the Atlantic Meridional Overturning Circulation (MOC). Would the baseline case that you have already run be a suitable comparison case? Why?
5. Set up a fully coupled case with the ocean model on a one-degree grid. What fully coupled compsets are available to you?
6. Set up and run a case in which you change ocean-model mixing coefficients. What are your options? What would you choose to do and why?
7. How would you set up a case to change the pressure formulation in the sea-ice model?
8. Your baseline experiment uses the ocean ecosystem model. How would you find out how to modify the ecosystem-model parameters? How would you configure an otherwise identical G case, except with the ecosystem model turned off? Would you expect the solutions from this case to be identical to your baseline case?